

Document made available under the Patent Cooperation Treaty (PCT)

International application number: PCT/CA04/002198

International filing date: 24 December 2004 (24.12.2004)

Document type: Certified copy of priority document

Document details: Country/Office: US
Number: 60/562,240
Filing date: 15 April 2004 (15.04.2004)

Date of receipt at the International Bureau: 22 April 2005 (22.04.2005)

Remark: Priority document submitted or transmitted to the International Bureau in compliance with Rule 17.1(a) or (b)




World Intellectual Property Organization (WIPO) - Geneva, Switzerland
Organisation Mondiale de la Propriété Intellectuelle (OMPI) - Genève, Suisse

PA 1300690

PCT/CA

04002198

CA04/2198



THE UNITED STATES OF AMERICA

TO ALL TO WHOM THESE PRESENTS SHALL COME:

UNITED STATES DEPARTMENT OF COMMERCE

United States Patent and Trademark Office


March 28, 2005

THIS IS TO CERTIFY THAT ANNEXED HERETO IS A TRUE COPY FROM THE RECORDS OF THE UNITED STATES PATENT AND TRADEMARK OFFICE OF THOSE PAPERS OF THE BELOW IDENTIFIED PATENT APPLICATION THAT MET THE REQUIREMENTS TO BE GRANTED A FILING DATE UNDER 35 USC 111.

APPLICATION NUMBER: 60/562,240

FILING DATE: April 15, 2004

**By Authority of the
COMMISSIONER OF PATENTS AND TRADEMARKS**



**M. SIAS
Certifying Officer**

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PROVISIONAL APPLICATION FOR PATENT COVER SHEET

This is a request for filing a PROVISIONAL APPLICATION FOR PATENT under 37 CFR 1.53(c).

Express Mail Label No. 13441 U.S. PTO
60/562240

041504

INVENTOR(S)					
Given Name (first and middle [if any])		Family Name or Surname		Residence (City and either State or Foreign Country)	
Lorne Helge		WHITEHEAD SEETZEN		Vancouver, CANADA Montreal, CANADA	
Additional inventors are being named on the <u>1</u> separately numbered sheets attached hereto					
TITLE OF THE INVENTION (500 characters max)					
Low Dynamic Range Image Enhancement for High Dynamic Range Systems					
Direct all correspondence to: CORRESPONDENCE ADDRESS					
<input type="checkbox"/> Customer Number:		<div style="border: 1px solid black; width: 250px; height: 25px;"></div>			
OR					
<input checked="" type="checkbox"/> Firm or Individual Name		The University of British Columbia - Industry Liaison Office			
Address		#103 -6190 Agronomy Road			
Address					
City	Vancouver	State	BC	Zip	V6T 1Z3
Country	CANADA	Telephone	6048228594	Fax	6048225998
ENCLOSED APPLICATION PARTS (check all that apply)					
<input checked="" type="checkbox"/> Specification Number of Pages <u>18</u>		<input type="checkbox"/> CD(s), Number _____			
<input type="checkbox"/> Drawing(s) Number of Sheets _____		<input type="checkbox"/> Other (specify) _____			
<input type="checkbox"/> Application Date Sheet. See 37 CFR 1.76					
METHOD OF PAYMENT OF FILING FEES FOR THIS PROVISIONAL APPLICATION FOR PATENT					
<input checked="" type="checkbox"/> Applicant claims small entity status. See 37 CFR 1.27.		FILING FEE Amount (\$) <div style="border: 1px solid black; width: 100px; height: 50px; margin: 0 auto; text-align: center; line-height: 50px;">80.00</div>			
<input type="checkbox"/> A check or money order is enclosed to cover the filing fees.					
<input type="checkbox"/> The Director is hereby authorized to charge filing fees or credit any overpayment to Deposit Account Number: _____					
<input checked="" type="checkbox"/> Payment by credit card. Form PTO-2038 is attached.					
The invention was made by an agency of the United States Government or under a contract with an agency of the United States Government.					
<input checked="" type="checkbox"/> No.					
<input type="checkbox"/> Yes, the name of the U.S. Government agency and the Government contract number are: _____					

[Page 1 of 2]

Respectfully submitted,

SIGNATURE

TYPED or PRINTED NAME ANNARI FAURIETELEPHONE 604-822-8594Date 14 April 2004

REGISTRATION NO. _____

(if appropriate)

Docket Number: UBC 04-028**USE ONLY FOR FILING A PROVISIONAL APPLICATION FOR PATENT**

This collection of information is required by 37 CFR 1.51. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 8 hours to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Mail Stop Provisional Application, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PROVISIONAL APPLICATION COVER SHEET
Additional Page

PTO/SB/16 (08-03)

Approved for use through 07/31/2006. OMB 0851-0032

U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

Docket Number UBC 04-028

INVENTOR(S)/APPLICANT(S)

Given Name (first and middle [if any])	Family or Surname	Residence (City and either State or Foreign Country)
Greg	WARD	Albany, CA

[Page 2 of 2]

Number 1 of 1

WARNING: Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.

Low Dynamic Range Image Enhancement for High Dynamic Range Systems

Introduction

Recent developments in the fields of image recording and display hardware have reached a point where emerging recording systems are capable of capturing the full luminance range found in real scenes and new display systems can show the same full luminance range or a close approximation of it. Such devices are called High Dynamic Range (HDR) cameras (or other image recording devices) and displays respectively.

Conventional cameras/displays capture/present only a very limited subset of the full luminance present in real scenes. These devices are called Low Dynamic Range (LDR). Most conventional LDR systems use 8-bits per colour channel (256 distinctly addressable steps within the full luminance range of the display). HDR devices use more than 8-bits per colour channel (usually 10, 12 or 16 bits for common configurations). This development creates a requirement for transition technology not only to interface the new high dynamic range image recording hardware with the corresponding displays but also for the configurations where one side of the interface is still a low dynamic range device. In the near future this is expected to be the most common kind of interface as high dynamic range recording and display technologies are only just beginning to enter the imaging market, which is dominated by conventional low dynamic range devices.

The interface of high dynamic range recording devices to low dynamic range displays has already been developed for the purpose of compressing computer generated high dynamic range renderings into low dynamic range image formats. This process is generally referred to as tone mapping. This patent application describes basic elements for the reverse process of converting low dynamic range image data into content for high dynamic range displays. Such technology is applicable for the conversion of legacy data (old image data, photographs, film, etc) and for new data that is not captured with a high dynamic range capture device.

Brief Description of Invention

It is an object of this invention to provide a method and algorithm that is useful for converting image data that presents in a low dynamic range (e.g. where the image data contains 8 bits per colour channel) into image data that presents in a higher dynamic range (e.g. where the image data presents in 10, 12, or 16 bits per colour channel). The method is not limited to a particular bit depth and can be applied to increase the dynamic range of any input range to any output range of an image (e.g. 6-bit to 24-bit, etc) In its broadest embodiment the algorithm performs the following steps:

1. A first temporary copy of the image is created.
2. The lines of pixels are individually scanned along each line of a first axis (e.g. along each horizontal line of pixels) until a first saturated pixel is detected.
3. The length of the area of saturation (i.e. the number of contiguous saturated pixels) is established.
4. A parabola is computed centered on the center of the saturated area. If the saturated area extends to the edge of the image then the vertex of the parabola is placed on the edge of the image along the axis being scanned.
5. Pixels with luminance values defined by the parabola replace the pixels of the saturated pixel line.
6. Scanning is continued along the pixel line until a next saturated pixel is detected.
7. Steps 1 to 6 are repeated as necessary until the whole image has been scanned, the resulting image is stored.
8. Steps 1 to 7 are repeated for lines of pixels along each line of a second axis (e.g. along each vertical line of pixels). This step may optionally be repeated along other axes (e.g. diagonal lines).
9. The images created in each iteration of step 7 are combined to produce a final image.

The algorithm may be optionally improved by setting the maximum luminance at the vertex of the parabolas calculated in step 4 above by taking into account the luminance gradient of the pixels to the left and right of each area of saturation.

Detailed Description of Invention

In the following description all methods are described step by step as applicable to single images. Most of these methods involve only a few multiplications per image pixel and can be implemented in real time. They can consequently be implemented in software as part of an image editor program, on the graphic card of a computer, or as part of a dedicated signal processing unit inside a display (e.g. a television (TV) or computer monitor). In particular, it would be possible to use the presented enhancement tree to extend the dynamic range of the standard NTSC TV signal as long as the TV unit has some moderate processing power, which is standard in the current generation of high-end TV systems. More complex versions of the basic enhancement tree could be used for time-insensitive procedures such as the extension of photographs.

All of the methods described below assume that the enhancement processes require no human interaction. The enhancement process could be significantly improved if a human user were to make certain decisions in the process, which is not out of the question where the technology is applied to the digital enhancement of photographs.

One method of extending an 8-bit image to a 16-bit high dynamic range image is the linear extension of the 8-bit range via multiplication of each 8-bit value by 256. This effectively stretches the 8-bit image data across the 16-bit range and introduces a greatly enhanced level of contrast in the image; however, this technique may also generate unrealistic artefacts. In real world applications only a very small portion of scenes have a very high brightness level and a linear extension would not reflect this. More important, a linear extension does not discriminate between areas of saturation (either dark or white) and areas close to, but not at, saturation.

Nevertheless, a linear extension can somewhat improve the perceived image quality. A first step of any of the image enhancement steps described below can be any combination of linear extension (scaling) and adding an offset (i.e. sliding the entire image range upwards or downwards in level). These are standard processes implemented in most normal image editing software tools.

The following method can compensate for clamping, which occurs when high dynamic range data is compressed into a low dynamic range environment and either the top or bottom end of the range is cut off. It is this kind of saturation that linear extension cannot correct. For the sake of simplicity, in the following description we will only consider the effects of the upper level boundary (clamping level) since it is more commonly exceeded in real scenes than the lower level boundary. All presented methods work equally well for extension at the lower level boundary. In the following description it is assumed that image levels are in some fashion related to the illuminance values of the real scene that was the source of the image.

Ideally, the saturated areas are reconstructed by predicting how the image levels vary within the area of saturation. Unfortunately, there are no direct indicators for this variation and, in the absence of an intelligent (e.g. Human) decision maker, it is impossible to tell whether a particular area of saturation corresponded to a very bright source of light or just a shallow extension of the surrounding illuminance values that just happened to barely exceed the upper illuminance boundary. Without human interactions this method instead relies on the two known properties of a saturated area: its size and the behaviour of the image pixels close to the area of saturation. Using these two properties the method can make an attempt to reconstruct the illuminance of the scene within the area of saturation. The core of the invention described in this application is the use of known image characteristics of an area of saturation to make an approximation of the content of the saturated area and thereby fill in this content. This process enhances low dynamic range images by adding highlights in areas of upper level clamping and shadows in areas of lower level clamping. The following describes an example process using a specific enhancement technique (linear enhancement, offset, vertex extension of saturated areas based on saturation size and behaviour of the image around the saturation).

Starting with the size of the area, it is reasonable to assume that, in general, larger saturated areas correspond to brighter areas in real scenes when compared to smaller saturated areas. There are obviously exceptions where small areas situated within a darker surrounding are nevertheless very bright (e.g. street lamps at night, etc) and this must be kept in mind as the various enhancement steps to avoid obvious artefacts are designed. At least to first order though let us proceed with a straightforward extension of each saturated area based on its size. The simplest way to do so is by using a parabola centered on the area with a vertex luminance proportional to the size of the area.

$$L(x) = \frac{1}{4} \frac{x^2 - 2vx + v^2 + 4L_v k}{k}$$

Equation 1: Parabola equation for saturation extension. (v = vertex pixel, L_v = vertex luminance, k = steepness parameter, x = current pixel)

In order to implement such a parabola extension we can break the computation into three processing steps per image frame. In the first step the system scans through each horizontal line of pixels until a saturated pixel is encountered (i.e. a pixel with an 8-bit value of 255 or the new maximum greyscale value after a limited linear extension). Once the first saturated pixel is encountered, the system establishes the horizontal length of the area of saturation (i.e. the number of pixels until a non-saturated pixel is reached). Based on the length of the line of saturation the appropriate parabola is computed and the corresponding new pixel values replace the saturated pixel line. If the saturation line extends towards the edge of the image, the vertex of the parabola is placed onto the edge.




Figure 1: Schematic representation of basic enhancement process. Height represents abstract luminance values. Top left: Original 8-bit image featuring a saturated plateau (light blue region). Top right: Result of horizontal extension. Bottom left: Result of vertical extension. Bottom right: Final image (average of horizontal and vertical results).

The easiest way to implement this process is by indexing the saturated pixel line such that the pixel receives integer indices starting with the index 0 for the first encountered pixel. In that case, the vertex pixel index is given by half the maximum index for the saturation line. We can arbitrarily define the vertex luminance L_v and our choice defines the parameter k if we impose the boundary condition that the parabola needs to match up with the surrounding image at the first pixel of the saturated area with luminance L_0 .

$$k = \frac{1}{4} \frac{v^2}{L_0 - L_v}$$

Equation 2: Steepness parameter

The same procedure is repeated along the vertical pixel lines on a second temporary copy of the image. Once both processes have been completed, the enhanced image is generated from the average of the horizontal and vertically enhanced temporary copies.

For a many of images, this simple extension will vastly improve the appearance of previously saturated areas ; however, several potential artefacts may be encountered:

- The scaling constant k must be chosen with care and should be set on the conservative side. If too high a value of k is chosen then many artefacts will be seen where relatively dim areas (which are brighter than the upper illuminance bound) with large spatial sizes are portrayed as too bright. A conservative choice of k will largely eliminate these artefacts. The immediate trade-off to this conservatism is that small and very bright areas (i.e. street lamps, etc) are not sufficiently extended. Fortunately, this is a minor artefact as the human visual system is largely unable to assess the brightness of small areas accurately.
- Saturated areas extending to the edge of the image will be presented as being brightest at their edges, which is often not appropriate. Unfortunately, it is hard to imagine an indicator that would allow us to place the vertex of the parabola anywhere but on the edge since it is equally likely that the actual brightest area of the saturation is within or outside of the image. As long as k is conservative, this is not expected to lead to very significant artefacts.

The enhancement method is improved by taking the behaviour of the image close to the area of saturation into account. In real scenes, very high contrast boundaries are rare and as a result, most saturated areas are likely to be surrounded by a luminance gradient. It is reasonable to assume that the slope of this gradient generally extends into the saturated area. If the luminance gradient around the area of saturation is shallow then in general the area of saturation itself will be shallow, even if large. A steep gradient in turn will generally indicate a very bright area. This is true even if the bright object is sharply separated from the dark background because the light coming from the object will scatter inside the camera lens. The resulting halo will create an effective luminance gradient around the bright object.

The vertex luminance L_v can thus be computed by considering the average of the luminance slope before and after the line of saturation (lines are used rather than areas because the enhancement is conducted in two stages, along horizontal and vertical lines as before). It is reasonable to extend the analysis of the surrounding luminance gradient to the length of the saturation and thus the slope on each side of the saturation can be estimated as the average of all line slopes in that region

$$\frac{dL}{dx}_{avg} = \sum_{i=-v}^0 \frac{L_0 - L_i}{i}$$

Equation 2: Luminance slope to the left of saturation

Using these slope averages for the pixel to the left and right of the saturation, the vertex luminance can be determined by simple linear extrapolation. There might be more sophisticated models to determine the vertex luminance but this method is easily implemented and can be executed very quickly during image editing procedures.

$$L_v = \frac{1}{2} \nu \left[\left(\sum_{i=-\nu}^0 \frac{L_0 - L_i}{i} \right) + \left(\sum_{i=2\nu}^{3\nu} \frac{L_{2\nu} - L_i}{i} \right) \right] + L_0$$

Equation 3: Vertex luminance based on luminance gradient average surrounding the area of saturation

There are exceptions where this model fails. A good example of such an exception is a lamp on a white ceiling. Here the gradient will be very shallow but in the real scene the ceiling lamp will be approximately ten times brighter than the ceiling itself. Fortunately, this is the kind of conservative artefact we encountered before where the result of the enhancement is not bright enough but at least not excessively bright. By using the line slope average rather than the change in slope as we approach the saturation, we also avoid the possibility of a radical over-enhancement due to a sudden high contrast boundary (e.g. a checkerboard pattern where the white fields are saturated). Such a boundary would have a very steep slope and the change in slope would be dramatic.

To finish the enhancement algorithm, rules for the common exceptions must be defined:

If the separation between two saturated lines is smaller than the sum of the vertex positions of each line then the non-saturated part between the two saturations is split in half for the purpose of establishing the luminance gradient slope.

The maximum enhancement of a saturated area cannot exceed the luminance of the display system. Usually this cannot happen unless a very dramatic linear extension has been used prior to the saturation enhancement. In that case, the system will choose the highest possible vertex luminance and model the parabola accordingly (as opposed to creating a parabola with a vertex luminance exceeding the luminance range of the display system and thus creating saturation again).

Overall, the parabola enhancement allows an easily implemented method to improve the appearance of low dynamic range images significantly. A conservative choice of parameters ensures that the method does not create any obvious over-enhancement and thus the only artefacts introduced by the method are under-enhancement of some areas. The position of the vertex in the centre of saturation (or on the edge) remains an estimation and an improved system may base the vertex position on the difference in the luminance gradients approaching the saturation from the left and right. It is doubtful whether such a modification would significantly enhance the appearance of the image any further given that our visual system cannot accurately identify the brightest point within an area of brightness.

Example Enhancement

To illustrate the process tree for a normal image enhancement we can use the image of Stanford Memorial Church by Paul Debevec. Since the image is available in both a high dynamic range format and a tone-mapped 8-bit format, we can use it to compare the results of our enhancement with the actual desired luminance values.

Unfortunately, paper has only a very limited dynamic range and as a result the visualisation of the extended 16-bit image data requires some method of compressing the extended 16-bit images back into the approximate 8-bit range that can be shown in print. For all the following 16-bit images, this compression has been performed by directly dividing all image data by 255 to achieve a linear compression into the 8-bit range. Obviously, such a compression cannot be performed for the input 8-bit image, as it would compress the entire image into a single black layer. For the comparison of the 8-bit and extended 16-bit images, it is important to keep in mind that this compression mismatch between originally 8-bit and 16-bit images will make the 16-bit images appear darker on paper. This is of course not the case if the images are actually shown on a high dynamic range display system.

In order to allow for a better visualisation of the detailed processes during the extension, each image is accompanied by a graph showing the pixel values for the 470th line on the image (see Figure 3). To simplify the visualisation all of the images are presented in greyscale only. Extending a colour image follows the same steps and can be done with the same method as long as each channel is treated separately.




Figure 2: 8-bit input image




Figure 3: 470th row pixel values in 8-bit image. A small saturation can be observed at pixel 392 to 395 followed by a larger saturation over 20 pixels. The saturation corresponds to the third of the lower windows in the church image.

To start the extension we need to make a choice regarding the extension of the 8-bit image data (see Figure 2) into the 16-bit range. Clearly, a linear extension with a scale of 256 is not suitable if we want to continue the enhancement with saturation extension. A scale of 100 and an offset of 1000 will give us sufficient range below and above the image range to treat the saturations (see Figure 4). This will map the image range of 0 to 255 onto the 1000 to 26600 range of the 16-bit range. The offset of 1000 is sufficient for this particular image because there are few small areas of dark saturation in the image, all of which are shallow so that 1000 steps below the lowest image value will be more than sufficient.




Figure 4: Original image data multiplied by 100 and offset by 1000

The image is now ready for the extension algorithm (see Figure 5). The horizontal extension scans through the rows of the image as previously described and replaces all saturated parts with the appropriate parabola extension. On the 470th line, we can focus on the pixel range of 350 to 450 to get a closer look at the extension. The smaller four pixel wide saturation has only been extended by a very small amount because of the saturation size restriction imposed on the parabola. The twenty pixel wide saturation has been extended quite dramatically due to its size and the steepness of the image data leading up to the saturated pixel.




Figure 5: Horizontal extension of 470th row

In the actual image seen in Figure 6 the effect of the horizontal extension is clearly visible in the horizontal streaks in the top and rightmost bottom window. The two other windows as well as some of the reflections on the roof-beams also feature minor parabola extensions. Some of the dark regions in the ceiling area and in the left arch feature extensions of dark areas.




Figure 6: Horizontal extension of Stanford Memorial Church image

The vertical extension causes similar streaking in the vertical direction with different steepness and vertex positions of the parabolas. In the bottom right window for example the horizontal extensions are approximately 20 pixels across while the vertical extensions

BEST AVAILABLE COPY

are between 20 to 40 pixels across due to the geometry of the window. On the 470th row we notice this discrepancy in the fine structure of the vertical extensions of which we only see a slice on the 470th row (see Figure 7).




Figure 7: Comparison of horizontal and vertical parabola effects on 470th row. The oscillating height of the vertical extension is a result of the difference in length of the vertical columns of saturation intersecting with the 470th row in the window area.




Figure 8: Vertical extension of Stanford Memorial Church image

Finally, the two extension images can be superimposed (Figure 10) to obtain the final image (Figure 10). It is important to remember that all the images in print have been compressed into the 8-bit range. The still visible streaking in the images, in particular the top window is very noticeable in the compressed 8-bit version; however, this is almost

BEST AVAILABLE COPY

impossible to detect by the human eye if the image is shown on a high dynamic range display system. The streaking occurs almost at the top of the luminance range of the image and the maximum difference between the streaks in the bottom window is less than 6000 steps in the 470th row. The absolute image values near this maximum streak size are approximately 50,000 so the streak constitutes approximately a 10% effect. On a high dynamic range display system the 50,000th image value would correspond to a luminance of approximately 8,000cd/m². At that luminance level the contrast sensitivity of the eye under ideal conditions is slightly less than 10% and in a complex scene this is reduced much further. As the result, it is very unlikely that even the most significant streaks will be noticed by the viewer, especially if the image contains attention-drawing elements in a much lower luminance regime.




Figure 9: Final extension of 470th row. Averaging the horizontal and vertical results reduces the appearance of streaks.




Figure 10: Final result of the extension algorithm applied to the 8-bit compressed Stanford Memorial Church image.

The Stanford Memorial Church image is a very challenging scene for the parabola extension method because of the fine structure found in the saturated areas. The top

BEST AVAILABLE COPY

window is separated into segments by the metal window frame, which is occasionally broken on the photograph. This increases the length of some saturated strips significantly by connecting two saturated segments and the long vertical stripes found in the top windows after the extension are a direct result of this. The bottom windows are equally challenging since all three windows are made out of stained glass, which result in a very uneven pattern for the saturated areas. Despite these obstacles, the parabola extension yields a significantly improved final image with a minimum of visible artefacts. The same algorithm applied to a much more conventional image yields the same improvements without the streaking issue.





Figure 11: Comparison of original and extended office scene. The original image to the left has been linearly extended by a factor of 100 prior to re-compression into 8-bit to make the comparison easier. Clearly, the appearance of the ceiling lights and the desk lamp has improved during the extension but also notice the more subtle effect under the hanging cabinet.

As previously mentioned, the algorithm operates on colour images with the same results. The simplest modification to allow for colour image processing is to treat each colour channel as a separate image until they are recombined after the extension. The only additional artefact that could be introduced by this method for some colour images are slight discolouration of grey areas if they are not actually grey (i.e. equal red, green and blue values) but slightly biased towards one of the colour channels. In the colour version of the office scene such a discolouration effect is visible at the edge of the desk lamp where a greenish hue is noticeable.



BEST AVAILABLE COPY

Figure 12: Comparison of the colour images of the office scene (original to the left, extended to the right)

Summary

The method described in this application allows for extension of low dynamic range image data into high dynamic range data. This is done by appropriate scaling followed by enhancement of saturated or clamped areas of the image. These regions are extended according to the size of the area and the behaviour of the image close to the area of saturation. A basic algorithm for this approach has been illustrated. There are many different combinations or variations of this algorithm. Examples include:

- Different size dependency
- Different dependency on the surrounding pixel
- Treatment of colour images as luminance maps which are later multiplied with the colour channels
- More than horizontal and vertical line passes (e.g. diagonal passes, etc)